arXiv:1407.7220v2 [math.OC] 23 Sep 2014

A Parallel Method for Large Scale Convex Regression Problems

NECDET S. AYBAT[1] and ZI WANG[2]

***Abstract*— Convex regression (CR) problem deals with fitting a convex function to a finite number of observations. It has many applications in various disciplines, such as statistics, economics, operations research, and electrical engineering. Computing the least squares (LS) estimator via solving a quadratic program (QP) is the most common technique to fit a piecewise-linear convex function to the observed data. Since the number of constraints in the QP formulation increases quadratically in N, the number of observed data points, computing the LS estimator is not practical using interior point methods when N is very large. The first-order method proposed in this paper carefully manages the memory usage through parallelization, and efficiently solves large-scale instances of CR.**

I. Introduction

Convex regression (CR) problem is concerned with fitting a convex function to a finite number of observations. In particular, suppose that we are given N observations $\{(\boldsymbol{x}_\ell, \bar{y}_\ell)\}_{\ell=1}^N \subset \mathbb{R}^n \times \mathbb{R}$ such that

$$\bar{y}_\ell = f_0(\boldsymbol{x}_\ell) + \varepsilon_\ell, \quad \ell = 1, \dots, N, \tag{1}$$

where $f_0 : \mathbb{R}^n \to \mathbb{R}$ is convex, ε_ℓ is a random variable with $E[\varepsilon_\ell] = 0$ for all ℓ. The objective is to estimate the convex function f_0 from the observed data points. CR has many applications in various disciplines, such as statistics, economics, operations research, and electrical engineering. M. Mousavi [1] employed CR to estimate the value function under infinite-horizon discounted rewards for Markov chains, which naturally arises in various control problems. In economics field, CR is used for approximating consumers' concave utility functions from empirical data [2]. Moreover, in queueing network context, for models where the expectation of performance measure is convex in model parameters -see [3], using Monte Carlo methods to compute the expectation give rise to CR problem [4].

The most well-known method for CR is the least squares (LS) problem,

$$\hat{f}_N = \operatorname*{arg\,min}_{f \in \mathcal{C}} \sum_{\ell=1}^N \big(f(\boldsymbol{x}_\ell) - \bar{y}_\ell\big)^2, \tag{2}$$

where $\mathcal{C} := \{f : \mathbb{R}^n \to \mathbb{R} \text{ such that } f \text{ is convex}\}$. This infinite dimensional problem is equivalent to a finite dimensional quadratic problem (QP),

$$\begin{aligned} \min_{y_\ell \in \mathbb{R},\ \boldsymbol{\xi}_\ell \in \mathbb{R}^n} & \sum_{\ell=1}^N \big|y_\ell - \bar{y}_\ell\big|^2 \\ \text{s.t.} \quad & y_{\ell_1} \geq y_{\ell_2} + {\xi_{\ell_2}}^T(\boldsymbol{x}_{\ell_1} - \boldsymbol{x}_{\ell_2}) \quad 1 \leq \ell_1 \neq \ell_2 \leq N. \end{aligned} \tag{3}$$

[1]Industrial and Manufacturing Engineering Dept., Penn State University, University Park, PA 16802, USA. email: nsa10@psu.edu

[2]Industrial and Manufacturing Engineering Dept., Penn State University, University Park, PA 16802, USA. email: zxw121@psu.edu

Indeed, let $\{(y_\ell^*, \xi_\ell^*)\}_{\ell=1}^N$ be an optimal solution to (3), it is easy to show that when $N \geq n+1$, $\{y_\ell^*\}_{\ell=1}^N$ is unique, $\hat{f}_N(\boldsymbol{x}_\ell) = y_\ell^*$ and $\xi_\ell^* \in \partial \hat{f}_N(\boldsymbol{x}_\ell)$ for all ℓ, where ∂ denotes the subdifferential. Moreover, $\hat{f}_N \to f_0$ almost surely is shown in [4]; and the convergence rate is established in [5] for one-dimensional case, i.e. $n = 1$. LS estimator has some significant advantages over many other estimators proposed in the literature for CR. First, LS estimator is a non-parametric regression method as discussed in [6], which does not require any tuning parameters and avoids the issue of selecting an appropriate estimation structure. On the other hand, as discussed in [1], methods proposed by Hannah and Dunson [7], [8], are semi-parametric and require adjusting several parameters before fitting a convex function. Second, LS estimator can be computed by solving the QP in (3). Therefore, at least in theory, it can be solved very efficiently using interior point methods (IPM). However, a major drawback of LS estimator in practice is that the number of shape constraints in (3) is $\mathcal{O}(N^2)$. Consequently, the problem quickly becomes massive even for moderate number of observations: the complexity of each factorization step in IPM is $\mathcal{O}(N^3(n+1)^3)$, and the memory requirement of IPM is $\mathcal{O}\big(N^2(n+1)^2\big)$ assuming Cholesky factors are stored - see [9], [10].

In this paper, we develop a methodology for parallel computing the LS estimator on huge-scale CR problems. The proposed method carefully manages the memory usage through parallelization, and efficiently solves large-scale instances of (3). Indeed, by regularizing the objective in (3), we ensure the feasibility of primal iterates in the limit, and Lipchitz continuity of gradient of the dual function. These properties lead to the main result, Theorem 2, which provides error bounds on the distance between the LS estimator and the optimal solution to the regularized problem. In the rest of the paper, after examining the dual decomposition for large-scale CR instances, we briefly discuss a first-order augmented Lagrangian method for solving QP subproblems. Finally, we conclude with a number of numerical examples.

II. Methodology

Assume that $\{\varepsilon_\ell\}_{\ell=1}^N$ is uniformly bounded by some $B_\varepsilon > 0$, $f_0 : \mathbb{R}^n \to \mathbb{R} \cup \{+\infty\}$ is a convex function, and $\{\boldsymbol{x}_\ell\}_{\ell=1}^N$ is a set of independent and identically distributed (i.i.d.) random vectors in $\mathbb{R}^n$ having a common continuous distribution supported on the n-dimensional hypercube $\mathcal{H} := [-B_x, B_x]^n \subset \mathbf{ri}\,\mathbf{dom}(f_0)$ for some $B_x > 0$, where $\mathbf{ri}$ denotes the relative interior.

Consider (3) in the following compact form,

$$\begin{aligned} \min_{\boldsymbol{y}\in\mathbb{R}^N,\ \boldsymbol{\xi}\in\mathbb{R}^{Nn}} \quad & \tfrac{1}{2}\left\|\boldsymbol{y}-\bar{\boldsymbol{y}}\right\|_2^2 \\ \text{s.t.} \quad & A_1\ \boldsymbol{y}+A_2\ \boldsymbol{\xi}\geq 0, \end{aligned} \tag{4}$$

where $A_1 \in \mathbb{R}^{N(N-1)\times N}$ and $A_2 \in \mathbb{R}^{N(N-1)\times Nn}$ are the matrices corresponding to the constraints in (3). Let

$$(\boldsymbol{y}^*,\boldsymbol{\xi}^*) := \underset{\boldsymbol{y},\ \boldsymbol{\xi}}{\arg\min}\left\{\tfrac{1}{2}\left\|\boldsymbol{y}\right\|_2^2+\tfrac{1}{2}\left\|\boldsymbol{\xi}\right\|_2^2 :\ (\boldsymbol{y},\boldsymbol{\xi})\in\chi^*\right\}, \tag{5}$$

where χ^* denotes the set of optimal solutions to (4). Let $B_\xi := \|\boldsymbol{\xi}^*\|_2$. Moreover, since (4) is a convex QP, strong duality holds, and an optimal dual solution $\boldsymbol{\theta}^* \in \mathbb{R}^{N(N-1)}$ exists. Let $B_\theta > 0$ such that $\|\boldsymbol{\theta}^*\|_\infty \leq B_\theta$ for some optimal dual. The complexity result of the proposed method will be provided in terms of constants B_ξ and B_θ.

A. Separability

To reduce *curse of dimensionality* and develop a first-order parallel algorithm that can solve (4) for large N, we use dual decomposition to induce separability. To this aim, we partition N observations into K subsets. Let $\{\mathcal{C}_i\}_{i=1}^K$ denote the collection of indices such that $|\mathcal{C}_i| \geq n+1$ for all i. To simplify the notation, let $N = K\bar{N}$ for some $\bar{N} > n+1$,

$$\mathcal{C}_i := \left\{(i-1)\bar{N}+1,\ (i-1)\bar{N}+2,\ldots,\ i\bar{N}\right\}$$

for $1 \leq i \leq K$. Throughout the paper, $\boldsymbol{y}_i \in \mathbb{R}^{\bar{N}}$ and $\boldsymbol{\xi}_i \in \mathbb{R}^{\bar{N}n}$ denote the sub-vectors of $\boldsymbol{y} \in \mathbb{R}^N$ and $\boldsymbol{\xi} \in \mathbb{R}^{Nn}$ corresponding to indices in $\mathcal{C}_i$, respectively.

For every ordered pair (ℓ_1,ℓ_2) such that $1 \leq \ell_1 \neq \ell_2 \leq N$, there corresponds a constraint in (3) represented by a row in the matrices A_1 and A_2 of formulation (4). By dualizing all the constraints in (3) corresponding to $\ell_1 \neq \ell_2$ such that they belong to different sets in the partition, i.e. $\ell_1 \in \mathcal{C}_i$, $\ell_2 \in \mathcal{C}_j$ and $i \neq j$, we form the partial Lagrangian,

$$\begin{aligned} \mathcal{L}\left(\boldsymbol{y},\boldsymbol{\xi},\boldsymbol{\theta}\right) :=& \frac{1}{2}\sum_{i=1}^K \left\|\boldsymbol{y}_i-\bar{\boldsymbol{y}}_i\right\|_2^2 \\ & - \sum_{1\leq i\neq j\leq K}\sum \boldsymbol{\theta}_{ij}{}^{\mathsf{T}}\left(A_1^{ij}\begin{bmatrix}\boldsymbol{y}_i\\ \boldsymbol{y}_j\end{bmatrix}+A_2^{ij}\begin{bmatrix}\boldsymbol{\xi}_i\\ \boldsymbol{\xi}_j\end{bmatrix}\right), \end{aligned}$$

which will lead to the following partial dual function

$$\begin{aligned} g(\boldsymbol{\theta}) := \min_{\boldsymbol{y}\in\mathbb{R}^N,\ \boldsymbol{\xi}\in\mathbb{R}^{Nn}} \quad & \mathcal{L}\left(\boldsymbol{y},\boldsymbol{\xi},\boldsymbol{\theta}\right) \\ \text{s.t.} \quad & A_1^{ii}\ \boldsymbol{y}_i + A_2^{ii}\ \boldsymbol{\xi}_i \geq 0,\ i=1,\ldots,K, \end{aligned} \tag{6}$$

where for each $i \in \{1,\ldots,K\}$, A_1^{ii} and A_2^{ii} are formed by rows of A_1 and A_2, respectively, corresponding to all (ℓ_1,ℓ_2) such that $\ell_1 \neq \ell_2$ and $\ell_1,\ell_2 \in \mathcal{C}_i$; similarly, for each (i,j) such that $1 \leq i \neq j \leq K$, ${A_1}^{ij}$ and ${A_2}^{ij}$ contain the rows of A_1 and A_2, respectively, corresponding to $\{(\ell_1,\ell_2) :\ \ell_1 \in \mathcal{C}_i,\ \ell_2 \in \mathcal{C}_j\}$, and $\boldsymbol{\theta}_{ij} \in \mathbb{R}^{\bar{N}^2}$ denotes the associated dual variables. $\boldsymbol{\theta}$ denotes the vector formed by vertically concatenating $\boldsymbol{\theta}_{ij}$ for $1 \leq i \neq j \leq K$.

Note that partial Lagrangian $\mathcal{L}$ is separable and can be written as $\mathcal{L}\left(\boldsymbol{y},\boldsymbol{\xi},\boldsymbol{\theta}\right) = \sum_{i=1}^K \mathcal{L}_i\left(\boldsymbol{y}_i,\boldsymbol{\xi}_i,\boldsymbol{\theta}\right)$ for some very simple quadratic functions $\mathcal{L}_i$. Thanks to the separability of $\mathcal{L}$, computing the partial dual function $g(\boldsymbol{\theta})$, given in (6), is equivalent to solving K quadratic subproblems of the form:

$$\begin{aligned} \min_{\boldsymbol{y}_i\in\mathbb{R}^{\bar{N}},\ \boldsymbol{\xi}_i\in\mathbb{R}^{\bar{N}n}} \quad & \mathcal{L}_i(\boldsymbol{y}_i,\boldsymbol{\xi}_i,\boldsymbol{\theta}) \\ \text{s.t.} \quad & A_1^{ii}\ \boldsymbol{y}_i + A_2^{ii}\ \boldsymbol{\xi}_i \geq 0, \end{aligned} \tag{7}$$

for $1 \leq i \leq K$. Given the dual variables $\boldsymbol{\theta}$, since all K subproblems can be computed in parallel, one can take advantage of the computing power of multi-core processors. In the rest of the paper, we discuss how to compute a solution to (3) via solving the dual problem: $\max\{g(\boldsymbol{\theta}) :\ \boldsymbol{\theta} \geq 0\}$.

B. Projected Subgradient Method for Dual

One of the most well-known methods for solving the dual problem is the projected subgradient method. Let $\boldsymbol{\theta}_{ij}^0 = \mathbf{0}$ for all i, j such that $i \neq j$. Given the k-th dual iterate $\boldsymbol{\theta}^k$, $(\boldsymbol{y}^k,\boldsymbol{\xi}^k)$ denotes an optimal solution to the minimization problem in (6) when $\boldsymbol{\theta}$ is set to $\boldsymbol{\theta}^k$, and $\boldsymbol{\theta}_{ii}^*$ denotes an optimal dual associated with constraints $A_1^{ii}\ \boldsymbol{y}_i + A_2^{ii}\ \boldsymbol{\xi}_i \geq 0$ in (6). The next dual iterate $\boldsymbol{\theta}^{k+1}$ is computed as follows

$$\boldsymbol{\theta}_{ij}^{k+1} = \prod_{\mathcal{S}_{ij}}\left(\boldsymbol{\theta}_{ij}^k - t_k\left(A_1^{ij}\begin{bmatrix}\boldsymbol{y}_i^k\\ \boldsymbol{y}_j^k\end{bmatrix}+A_2^{ij}\begin{bmatrix}\boldsymbol{\xi}_i^k\\ \boldsymbol{\xi}_j^k\end{bmatrix}\right)\right), \tag{8}$$

where $\Pi_{\mathcal{S}_{ij}}(.)$ denotes the Euclidean projection on to

$$\mathcal{S}_{ij} = \left\{\boldsymbol{\theta}_{ij} \geq \mathbf{0} : {\boldsymbol{\theta}_{ij}}^{\mathsf{T}} A_2^{ij} + \begin{bmatrix}\boldsymbol{\theta}_{ii}^*{}^{\mathsf{T}} A_2^{ii} & \boldsymbol{\theta}_{jj}^*{}^{\mathsf{T}} A_2^{jj}\end{bmatrix} = \mathbf{0}\right\}.$$

Since $\mathcal{L}$ is linear in $\boldsymbol{\xi}$, $\mathbf{dom}\, g$ is non-trivial and is given by the Cartesian product of $\mathcal{S}_{ij}$'s. The projected subgradient method is guaranteed to converge in function value with a careful selection of step size sequence $\{t_k\}_{k=1}^\infty$, and it requires $\mathcal{O}(1/\epsilon^2)$ iterations to obtain an ϵ-optimal solution - see [11]. However, due to lack of strong convexity of the objective function in (4) (not in $\boldsymbol{\xi}$), even if the dual variables converge to an optimal dual solution, the primal feasibility cannot be guaranteed in the limit.

C. Tikhonov Regularization Approach

In order to ensure feasibility in the limit, which cannot be guaranteed by the subgradient method discussed above, we employ Tikhonov regularization, of which convergence properties were investigated in [12]. Given $\gamma > 0$, consider

$$\begin{aligned} (\boldsymbol{y}(\gamma),\boldsymbol{\xi}(\gamma)) = \underset{\boldsymbol{y},\ \boldsymbol{\xi}}{\arg\min} \quad & \tfrac{1}{2}\left\|\boldsymbol{y}-\bar{\boldsymbol{y}}\right\|_2^2 + \tfrac{\gamma}{2}\left\|\boldsymbol{\xi}\right\|_2^2 \\ \text{s.t.} \quad & A_1\boldsymbol{y} + A_2\boldsymbol{\xi} \geq 0. \end{aligned} \tag{9}$$

As γ decreases to zero from above, the minimizer $(\boldsymbol{y}(\gamma),\boldsymbol{\xi}(\gamma))$ converges to $(\boldsymbol{y}^*,\boldsymbol{\xi}^*)$ defined in (5).

Lemma 1: The minimizer of (9), $\boldsymbol{y}(\gamma)$ as a function of regularization parameter γ, is Hölder continuous on $[0,\infty)$,

$$\left\|\boldsymbol{y}(\gamma)-\boldsymbol{y}^*\right\|_2 \leq B_\xi\sqrt{\gamma}. \tag{10}$$

Proof: Let $(\boldsymbol{y}(\gamma),\boldsymbol{\xi}(\gamma))$ be the optimal solution to (9) and $(\boldsymbol{y}^*,\boldsymbol{\xi}^*)$ be defined as in (5). From the first-order optimality conditions of (9) and (4), we have

$$\begin{pmatrix}\boldsymbol{y}(\gamma)-\bar{\boldsymbol{y}}\\ \gamma\ \boldsymbol{\xi}(\gamma)\end{pmatrix}^{\mathsf{T}}\begin{pmatrix}\boldsymbol{y}^*-\boldsymbol{y}(\gamma)\\ \boldsymbol{\xi}^*-\boldsymbol{\xi}(\gamma)\end{pmatrix} \geq 0, \tag{11}$$

$$\begin{pmatrix} \boldsymbol{y}^* - \bar{\boldsymbol{y}} \\ 0 \end{pmatrix}^{\mathsf{T}} \begin{pmatrix} \boldsymbol{y}(\gamma) - \boldsymbol{y}^* \\ \boldsymbol{\xi}(\gamma) - \boldsymbol{\xi}^* \end{pmatrix} \geq 0. \tag{12}$$

Note that both $(\boldsymbol{y}(\gamma), \boldsymbol{\xi}(\gamma))$ and $(\boldsymbol{y}^*, \boldsymbol{\xi}^*)$ are feasible to (4) and (9). This implies $\|\boldsymbol{\xi}(\gamma)\|_2 \leq \|\boldsymbol{\xi}^*\|_2$. Summing up (11) and (12), and using $B_\xi = \|\boldsymbol{\xi}^*\|_2$, it follows that

$$\|\boldsymbol{y}(\gamma) - \boldsymbol{y}^*\|_2^2 \leq \gamma \boldsymbol{\xi}(\gamma)^{\mathsf{T}} \big(\boldsymbol{\xi}^* - \boldsymbol{\xi}(\gamma)\big) \leq \gamma B_\xi^2.$$

□

Since the objective function in (9) is strongly convex in both $\boldsymbol{y}$ and $\boldsymbol{\xi}$, Danskin's theorem (see [13]) implies that the Lagrangian dual function of (9) is differentiable; therefore, one can use gradient type methods to solve the corresponding dual problem. Moreover, strong convexity ensures that, one can solve the primal problem by solving the dual problem. Indeed, let $\boldsymbol{\theta}(\gamma)$ be an optimal solution to the dual problem of (9), we can recover $(\boldsymbol{y}(\gamma), \boldsymbol{\xi}(\gamma))$ by computing the primal minimizers in (6) when the dual is set to $\boldsymbol{\theta}(\gamma)$. The discussion above shows that achieving primal feasibility is not an issue provided that we can solve the dual of (9). This motivates the next section, where we briefly state a first-order algorithm that can efficiently solve the dual of (9).

D. Accelerated Proximal Gradient (APG) Algorithm

Let $\rho : \mathbb{R}^d \to \mathbb{R}$ be a concave function such that $\nabla\rho$ is Lipschitz continuous on $\mathbb{R}^d$ with constant L, and $\mathcal{Q} \subset \mathbb{R}^d$ be a compact convex set. The APG algorithm [14], [15] displayed in Figure 1 is based on Nesterov's accelerated gradient method [11], [16] and solves $\rho^* = \max\{\rho(\eta) : \eta \in \mathcal{Q}\}$. Corollary 3 in [15], and Theorem 4.4 in [14] show that for all $\ell \geq 1$ the error bound is given by

$$0 \leq \rho^* - \rho(\eta_\ell) \leq \frac{2L}{(\ell+1)^2} \|\eta_0 - \eta^*\|_2^2,$$

where η_0 is the initial APG iterate and $\eta^* \in \arg\min_{\eta \in \mathcal{Q}} \rho(\eta)$. Hence, using APG one can compute an δ-optimal solution within at most $\mathcal{O}(\sqrt{L/\delta})$ APG iterations.

Algorithm APG $\left(\ \eta_0\ \right)$
Iteration 0: Take $\eta_0^{(1)} = \eta_1^{(2)} = \eta_0, t_1 = 1$
Iteration ℓ: ($\ell \geq 1$) Compute
1) $\eta_\ell^{(1)} = \Pi_{\mathcal{Q}} \left(\eta_\ell^{(2)} + \frac{\nabla\rho(\eta_\ell^{(2)})}{L} \right)$
2) $t_{\ell+1} = (1 + \sqrt{1 + 4t_\ell^2})/2$
3) $\eta_{\ell+1}^{(2)} = \eta_\ell^{(1)} + \frac{t_\ell - 1}{t_{\ell+1}} \left(\eta_\ell^{(1)} - \eta_{\ell-1}^{(1)} \right)$

Fig. 1: Accelerated Proximal Gradient Algorithm

In this paper, we will use APG algorithm on a slightly different but equivalent problem to (9). Let A_3 and A_4 denote the matrices formed by vertically concatenating A_1^{ij} and A_2^{ij}, respectively, for $1 \leq i \neq j \leq K$; and define

$$C = \begin{bmatrix} A_3 & A_4 \\ I & 0 \end{bmatrix}, \tag{13}$$

where $I \in \mathbb{R}^{N \times N}$ identity matrix. For notational convenience, let $\boldsymbol{\eta}^{\mathsf{T}} = \begin{bmatrix} \boldsymbol{y}^{\mathsf{T}} & \boldsymbol{\xi}^{\mathsf{T}} \end{bmatrix}$, and consider

$$\min_{\boldsymbol{\eta} \in Q_1} \tfrac{1}{2} \|\boldsymbol{y} - \bar{\boldsymbol{y}}\|_2^2 + \tfrac{\gamma}{2} \|\boldsymbol{\xi}\|_2^2, \quad \text{s.t.} \quad C\,\boldsymbol{\eta} \geq 0, \tag{14}$$

where $Q_1 := \big\{(\boldsymbol{y}, \boldsymbol{\xi}) : A_1^{ii}\, \boldsymbol{y}_i + A_2^{ii}\, \boldsymbol{\xi}_i \geq 0, 1 \leq i \leq K\big\}$. Note that (14) is different from (9) only in constraints $\boldsymbol{y} \geq \mathbf{0}$. Via possibly shifting all the observations $\{\boldsymbol{y}_\ell\}_{\ell=1}^N$ up by a sufficiently large quantity, we can assume without loss of generality that $\boldsymbol{y}^*(\gamma) \geq \mathbf{0}$ under bounded error assumption, i.e. $|\varepsilon_\ell| \leq B_\varepsilon$ for all ℓ. Therefore, (9) and (14) are indeed equivalent problems. Consider the dual problem of (14),

$$\max_{\boldsymbol{\theta}} g_\gamma(\boldsymbol{\theta}) \quad \text{s.t.} \quad \boldsymbol{\theta} \in Q_2, \tag{15}$$

where $Q_2 := \big\{\boldsymbol{\theta} : \|\boldsymbol{\theta}\|_2 \leq B_\theta,\ \boldsymbol{\theta} \geq 0\big\}$, and

$$g_\gamma(\boldsymbol{\theta}) = \min_{(\boldsymbol{y}, \boldsymbol{\xi}) \in Q_1} \left\{ \tfrac{1}{2} \big\|\boldsymbol{y} - \bar{\boldsymbol{y}}\big\|_2^2 + \tfrac{\gamma}{2} \|\boldsymbol{\xi}\|_2^2 - \boldsymbol{\theta}^{\mathsf{T}} C \boldsymbol{\eta} \right\}. \tag{16}$$

Let $\boldsymbol{\eta}(\boldsymbol{\theta})$ be the minimizer in (16). Theorem 7.1 in [17] and Danskin's theorem imply that

$$\nabla g_\gamma(\boldsymbol{\theta}) = -C\boldsymbol{\eta}(\boldsymbol{\theta}) \tag{17}$$

is Lipschitz continuous with constant

$$L_g = \tfrac{1}{\gamma}\, \sigma_{\max}^2(C). \tag{18}$$

Parallel APG algorithm (P-APG), displayed in Fig. 2, is the customized version of APG algorithm in Fig. 1 to solve (15). Note that at each iteration computation in Step 1) can be done in parallel using K processors, each solving a smaller QP.

Algorithm P-APG $\left(\ \gamma\ \right)$
Iteration 0: Take $\boldsymbol{\theta}_0^{(1)} = \boldsymbol{\theta}_1^{(2)} = \mathbf{0}, t_1 = 1$
Iteration ℓ: ($\ell \geq 1$) Compute
1) $\boldsymbol{\eta}_\ell = \underset{(\boldsymbol{y}, \boldsymbol{\xi}) \in Q_1}{\arg\min} \left\{ \tfrac{1}{2} \big\|\boldsymbol{y} - \bar{\boldsymbol{y}}\big\|_2^2 + \tfrac{\gamma}{2} \|\boldsymbol{\xi}\|_2^2 - \left(\boldsymbol{\theta}_\ell^{(2)}\right)^{\mathsf{T}} C\boldsymbol{\eta} \right\}$
2) $\boldsymbol{\theta}_\ell{}^{(1)} = \Pi_{Q_2} \left(\boldsymbol{\theta}_\ell^{(2)} - \frac{1}{L_g} C\boldsymbol{\eta}_\ell \right)$
3) $t_{\ell+1} = (1 + \sqrt{1 + 4t_\ell^2})/2$
4) $\boldsymbol{\theta}_{\ell+1}^{(2)} = \boldsymbol{\theta}_\ell^{(1)} + \frac{t_\ell - 1}{t_{\ell+1}} \left(\boldsymbol{\theta}_\ell^{(1)} - \boldsymbol{\theta}_{\ell-1}^{(1)} \right)$

Fig. 2: Parallel APG Algorithm

Note that the iteration complexity of gradient ascent method on (15) is $\mathcal{O}(L_g/\delta) = \mathcal{O}(B_\theta^2(\gamma\delta)^{-1})$. On the other hand, P-APG in Fig. 2 can compute a δ-optimal solution to (15) within $\mathcal{O}(\sqrt{L_g/\delta})$ iterations. More precisely, (18) implies $\mathcal{O}(B_\theta(\gamma\delta)^{-1/2})$ complexity for P-APG on (15).

Let $\boldsymbol{\theta}_\delta$ be a δ-optimal solution to (15), and $(\boldsymbol{y}_\delta, \boldsymbol{\xi}_\delta)$ be the optimal solution to the minimization problem in (6) when $\boldsymbol{\theta}$ is set to $\boldsymbol{\theta}_\delta$. In Theorem 2, which is the main result of this paper, we establish an error bounds on suboptimality $\|\boldsymbol{y}_\delta - \boldsymbol{y}^*\|_2$, and on infeasibility $\|(A_1\boldsymbol{y}_\delta + A_2\boldsymbol{\xi}_\delta)_-\|_2$, where $(\boldsymbol{x})_- := \max\{-\boldsymbol{x}, \mathbf{0}\}$.

Theorem 2: Let $(\boldsymbol{y}(\gamma), \boldsymbol{\xi}(\gamma))$ and $\boldsymbol{\theta}^*$ denote the optimal solutions to (9) and (15), respectively. Let $\boldsymbol{\theta}_\delta$ be a δ-optimal solution to (15), and $(\boldsymbol{y}_\delta, \boldsymbol{\xi}_\delta)$ be the minimizer in (16) when $\boldsymbol{\theta}$ is set to $\boldsymbol{\theta}_\delta$. For all $\delta > 0$, the following bounds hold:

$$\|\boldsymbol{y}_\delta - \boldsymbol{y}^*\|_2 \leq B_\xi\sqrt{\gamma} + \sqrt{\tfrac{2\delta}{\gamma}}\sigma_{\max}(C), \tag{19}$$

$$\|(A_1\boldsymbol{y}_\delta + A_2\boldsymbol{\xi}_\delta)_-\|_2 \leq \sqrt{\tfrac{2\delta}{\gamma}}\sigma_{\max}(C). \tag{20}$$

Proof: Since g_γ is Lipschitz continuous with constant L_g given in (18), we have

$$\|\nabla g_\gamma(\boldsymbol{\theta}_1) - \nabla g_\gamma(\boldsymbol{\theta}_2)\|_2 \leq \frac{\sigma^2_{\max}(C)}{\gamma}\|\boldsymbol{\theta}_1 - \boldsymbol{\theta}_2\|_2.$$

Moreover, first order optimality conditions for (15) imply

$$-\langle \nabla g(\boldsymbol{\theta}^*),\quad \boldsymbol{\theta} - \boldsymbol{\theta}^*\rangle \geq 0, \qquad \forall \boldsymbol{\theta} \in Q_2. \tag{21}$$

From (2.1.7) in [11], it follows that

$$\begin{aligned}-\nabla g_\gamma(\boldsymbol{\theta}^*)^\mathsf{T}(\boldsymbol{\theta}_\delta - \boldsymbol{\theta}^*) + \frac{\gamma}{2\sigma^2_{\max}(C)}\|\nabla g_\gamma(\boldsymbol{\theta}_\delta) - \nabla g_\gamma(\boldsymbol{\theta}^*)\|_2^2 \\ \leq -g_\gamma(\boldsymbol{\theta}_\delta) + g(\boldsymbol{\theta}^*) \leq \delta.\end{aligned}$$

Using (13), (17) and (21), we have

$$\left\|\begin{bmatrix} A_3(\boldsymbol{y}(\gamma) - \boldsymbol{y}_\delta) + A_4(\boldsymbol{\xi}(\gamma) - \boldsymbol{\xi}_\delta) \\ \boldsymbol{y}(\gamma) - \boldsymbol{y}_\delta \end{bmatrix}\right\|_2^2 \leq \tfrac{2\delta}{\gamma}\sigma^2_{\max}(C). \tag{22}$$

Hence, together with (10), it implies (19). Moreover, since $\|\boldsymbol{x} - \boldsymbol{y}\|_2 \geq \|(\boldsymbol{x})_- - (\boldsymbol{y})_-\|_2$ for any $\boldsymbol{x}$ and $\boldsymbol{y}$, we also have

$$\|(A_3\boldsymbol{y}_\delta + A_4\boldsymbol{\xi}_\delta)_- - (A_3\boldsymbol{y}(\gamma) + A_4\boldsymbol{\xi}(\gamma))_-\|_2 \leq \sqrt{\tfrac{2\delta}{\gamma}}\sigma_{\max}(C).$$

Since $(\boldsymbol{y}(\gamma), \boldsymbol{\xi}(\gamma))$ is feasible to (9), and $(\boldsymbol{y}_\delta, \boldsymbol{\xi}_\delta) \in Q_1$, above inequality implies (20). □

Next, we prove an important technical lemma that will be used later in Theorem 4 to show that $\|\boldsymbol{\xi}_\delta - \boldsymbol{\xi}^*\|_2$ is small.

Lemma 3: Assuming that $\{\boldsymbol{x}_i\}_{i=1}^N$ are uniformly sampled at random from set $\phi = \{\boldsymbol{x} \in \mathbb{R}^n : \|\boldsymbol{x}\|_\infty \leq B_x\}$, the matrix A_4 in (13) has linearly independent columns (LIC).

Proof: Remember $A_4 \in \mathbb{R}^{N(N-\bar{N})\times Nn}$ denotes the matrix formed by vertically concatenating all A_2^{ij} for $1 \leq i \neq j \leq K$. Note that rows of $[A_1^{ij}\ \ A_2^{ij}]$ correspond to constraints $y_{\ell_1} - y_{\ell_2} + \boldsymbol{\xi}_{\ell_2}^\mathsf{T}(\boldsymbol{x}_{\ell_2} - \boldsymbol{x}_{\ell_1}) \geq 0$, where $\ell_1 \in \mathcal{C}_i$, $\ell_2 \in \mathcal{C}_j$ and $i \neq j$. For the sake of simplifying the discussion below, without loss of generality, we fix $i = 2$ and $j = 1$, and focus on the structure of A_2^{21}. Let $\hat{A}_2^{21}$ denote the submatrix of A_2^{21} formed by selecting the rows corresponding to $(\ell_1, \ell_2) \in \mathcal{C}_2 \times \mathcal{C}_1$ such that $\ell_1 = \bar{N} + 1$ and $\ell_2 \in \mathcal{C}_1$. Hence, we have

$$\hat{A}_2^{21} = \begin{bmatrix}\mathbf{X} & \mathbf{0}\end{bmatrix} \tag{23}$$

where $\mathbf{0}$ is the matrix of zeros and $\mathbf{X} \in \mathbb{R}^{\bar{N}\times\bar{N}n}$ such that

$$\mathbf{X} = \begin{bmatrix} \bar{\boldsymbol{x}}_1^\mathsf{T} & \mathbf{0}^\mathsf{T} & \mathbf{0}^\mathsf{T} & \dots & \mathbf{0}^\mathsf{T} \\ \mathbf{0}^\mathsf{T} & \bar{\boldsymbol{x}}_2^\mathsf{T} & \mathbf{0}^\mathsf{T} & \dots & \mathbf{0}^\mathsf{T} \\ \mathbf{0}^\mathsf{T} & \mathbf{0}^\mathsf{T} & \bar{\boldsymbol{x}}_3^\mathsf{T} & \dots & \mathbf{0}^\mathsf{T} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ \mathbf{0}^\mathsf{T} & \dots & \mathbf{0}^\mathsf{T} & \dots & \bar{\boldsymbol{x}}_{\bar{N}}^\mathsf{T} \end{bmatrix} \tag{24}$$

and $\bar{\boldsymbol{x}}_\ell := \boldsymbol{x}_\ell - \boldsymbol{x}_{\bar{N}+1}$ for $1 \leq \ell \leq \bar{N}$.

Fix $1 \leq j \leq K$. Note that for each $\ell \in \mathcal{C}_j$, there corresponds n columns in A_4; and the zero structure in (23) implies that each column of A_4 corresponding to $\mathcal{C}_j$ is linearly independent with $\bar{N}n$ columns in A_4 corresponding to $\mathcal{C}_k$ with probability 1 (w.p. 1) for all $k \neq j$. Moreover, when we focus on (24), we also see that any one of the n columns in A_4 corresponding to $\bar{\ell} \in \mathcal{C}_j$ is also linearly independent with n columns in A_4 corresponding to $\ell \in \mathcal{C}_j$ with probability 1 for all $\ell \neq \bar{\ell}$. Therefore, to show that A_4 has linearly independent columns, it is sufficient to show that for any given $1 \leq j \leq K$ and $\ell \in \mathcal{C}_j$, the corresponding n columns of A_4 are linearly independent w.p. 1.

Let $D \in \mathbb{R}^{N(N-\bar{N})\times n}$ be the submatrix of $A_4 \in \mathbb{R}^{N(N-\bar{N})\times Nn}$ corresponding to columns $\bar{\ell} \in \mathcal{C}_j$ for some $1 \leq j \leq K$; and $\boldsymbol{d}_{\ell_1\ell_2}^\mathsf{T}$ denote the row of D corresponding to (ℓ_1, ℓ_2) such that ℓ_1 and ℓ_2 belong to different sets in the partition. Clearly,

$$\boldsymbol{d}_{\ell_1\ell_2}^\mathsf{T} = \begin{cases} \left(\boldsymbol{x}_{\bar{\ell}} - \boldsymbol{x}_{\ell_1}\right)^\mathsf{T}, & \text{if } \ell_1 \notin \mathcal{C}_j \text{ and } \ell_2 = \bar{\ell}; \\ \mathbf{0}^\mathsf{T}, & \text{o.w.} \end{cases} \tag{25}$$

Without loss of generality, we fix $j > 1$ and consider $\bar{D} \in \mathbb{R}^{\bar{N}\times n}$ which denotes the submatrix of D corresponding to the rows $\boldsymbol{d}_{\ell_1\ell_2}^\mathsf{T}$ such that $\ell_1 \in \mathcal{C}_1$ and $\ell_2 = \bar{\ell} \in \mathcal{C}_j$. The following discussion is true for any C_k such that $k \neq j$, but setting $k = 1$ simplifies the notation in $\bar{D}$.

$$\bar{D} = \begin{pmatrix} \boldsymbol{x}_{\bar{\ell}}^\mathsf{T} - \boldsymbol{x}_1^\mathsf{T} \\ \vdots \\ \boldsymbol{x}_{\bar{\ell}}^\mathsf{T} - \boldsymbol{x}_\ell^\mathsf{T} \\ \vdots \\ \boldsymbol{x}_{\bar{\ell}}^\mathsf{T} - \boldsymbol{x}_{\bar{N}}^\mathsf{T} \end{pmatrix}$$

It suffices to show that $\bar{D}$ has LIC. Since $\bar{N} \geq n + 1$ and $\{\boldsymbol{x}_\ell\}_{\ell=1}^N$ is a set of i.i.d. random vectors in $\mathbb{R}^n$ having a common *continuous* distribution, it can be shown that there exists n linearly independent rows of $\bar{D}$ w.p. 1. Thus, A_4 has LIC. □

Theorem 4: There exists $K_1, K_2 > 0$ such that

$$\|\boldsymbol{\xi}_\delta - \boldsymbol{\xi}^*\|_2 \leq K_1\sqrt{\gamma} + K_2\sqrt{\tfrac{\delta}{\gamma}}.$$

Proof: Since $\boldsymbol{y}^*$ is the unique optimal solution to (4), (5) implies that $\boldsymbol{\xi}^* = \arg\min\{\|\boldsymbol{\xi}\|_2 :\ A_1\boldsymbol{y}^* + A_2\boldsymbol{\xi} \geq \mathbf{0}\}$. Similarly, (9) implies that $\boldsymbol{\xi}(\gamma) = \arg\min\{\|\boldsymbol{\xi}\|_2 :\ A_1\boldsymbol{y}(\gamma) + A_2\boldsymbol{\xi} \geq \mathbf{0}\}$. Hence, for $\boldsymbol{h}(\gamma) := A_1(\boldsymbol{y}^* - \boldsymbol{y}(\gamma))$,

$$\boldsymbol{\xi}(\gamma) = \arg\min\{\|\boldsymbol{\xi}\|_2 :\ A_1\boldsymbol{y}^* + A_2\boldsymbol{\xi} \geq \boldsymbol{h}(\gamma)\}. \tag{26}$$

Sensitivity of projection onto parametric polyhedral sets was studied in [18]. Using Theorem 2.1 in [18] and (10), we have

$$\|\boldsymbol{\xi}(\gamma) - \boldsymbol{\xi}^*\|_2 \leq K\|\boldsymbol{h}(\gamma)\|_2 \leq K\sigma_{\max}(A_1)B_\xi\sqrt{\gamma}, \tag{27}$$

for some $K > 0$. Moreover, (22) and Lemma 3 imply that

$$\|\boldsymbol{\xi}_\delta - \boldsymbol{\xi}(\gamma)\|_2 \leq \frac{\sqrt{\frac{2\delta}{\gamma}}\sigma_{\max}(C) + \sigma_{\max}(A_3)\|\boldsymbol{y}(\gamma) - \boldsymbol{y}_\delta\|_2}{\sigma_{\min}(A_4)}$$

Hence, $\|\boldsymbol{\xi}_\delta - \boldsymbol{\xi}(\gamma)\|_2 \leq \frac{(\sigma_{\max}(A_3)+1)\sigma_{\max}(C)}{\sigma_{\min}(A_4)}\sqrt{\frac{2\delta}{\gamma}}$. □

E. ALCC - An Augmented Lagrangian Method

Now, we first briefly state a first-order algorithm to directly solve (9). Let $\bar{B}_y$ and $\bar{B}_\xi$ be given such that $\boldsymbol{y}(\gamma) \in \mathcal{Q}_y := \{\boldsymbol{y} : \|\boldsymbol{y} - \bar{\boldsymbol{y}}\|_2 \leq \bar{B}_y\}$, and $\boldsymbol{\xi}(\gamma) \in \mathcal{Q}_\xi := \{\boldsymbol{\xi} : \|\boldsymbol{\xi}\|_2 \leq \bar{B}_\xi\}$. Such $\bar{B}_y$ and $\bar{B}_\xi$ can be found easily, if we are given

a feasible solution $(\hat{\boldsymbol{y}},\hat{\boldsymbol{\xi}})$, i.e. $A_1\hat{\boldsymbol{y}} + A_2\hat{\boldsymbol{\xi}} \geq \mathbf{0}$. Indeed, selecting $\bar{B}_y = \bar{B}$ and $\bar{B}_\xi = B/\sqrt{\gamma}$ works, where $\bar{B} := \sqrt{\|\hat{\boldsymbol{y}} - \bar{\boldsymbol{y}}\|_2^2 + \gamma\|\hat{\boldsymbol{\xi}}\|_2^2}$. ALCC [19] computes a solution to (9) by inexactly solving a sequence of subproblems:

$$P_k^* := \min\{P_k(\boldsymbol{y},\boldsymbol{\xi}) : \ \boldsymbol{y} \in \mathcal{Q}_y, \boldsymbol{\xi} \in \mathcal{Q}_\xi\}, \qquad (28)$$
$$P_k(\boldsymbol{y},\boldsymbol{\xi}) := \tfrac{1}{2\mu_k}\|\boldsymbol{y} - \bar{\boldsymbol{y}}\|_2^2 + \tfrac{\gamma}{2\mu_k}\|\boldsymbol{\xi}\|_2^2 + h_k(\boldsymbol{y},\boldsymbol{\xi}),$$

where $h_k(\boldsymbol{y},\boldsymbol{\xi}) := \frac{1}{2}\| (A_1\boldsymbol{y} + A_2\boldsymbol{\xi} - \boldsymbol{\theta}_k)_- \|_2^2$, and $\{\boldsymbol{\theta}_k\}$ sequence is defined in Fig. 3. For $k \geq 1$, $h_k(\boldsymbol{y},\boldsymbol{\xi})$ is convex in $\boldsymbol{y}$ and $\boldsymbol{\xi}$ -see Lemma 3.1 in [19]. Moreover,

$$\nabla_{\boldsymbol{y}} h_k(\boldsymbol{y},\boldsymbol{\xi}) = -{A_1}^\mathsf{T}(A_1\boldsymbol{y} + A_2\boldsymbol{\xi} - \boldsymbol{\theta}_k)_-,$$
$$\nabla_{\boldsymbol{\xi}} h_k(\boldsymbol{y},\boldsymbol{\xi}) = -{A_2}^\mathsf{T}(A_1\boldsymbol{y} + A_2\boldsymbol{\xi} - \boldsymbol{\theta}_k)_-.$$

In addition, $\nabla_{\boldsymbol{y}} h_k(\boldsymbol{y},\boldsymbol{\xi})$ is Lipschitz continuous in $\boldsymbol{y}$ for all fixed $\boldsymbol{\xi}$ with constant $\sigma_{\max}^2(A_1)$, and $\nabla_{\boldsymbol{\xi}} h_k(\boldsymbol{y},\boldsymbol{\xi})$ is Lipschitz continuous in $\boldsymbol{\xi}$ for all fixed $\boldsymbol{y}$ with constant $\sigma_{\max}^2(A_2)$. Hence, $\nabla_{\boldsymbol{y}} P_k(\boldsymbol{y},\boldsymbol{\xi})$ is Lipschitz continuous in $\boldsymbol{y}$ for all fixed $\boldsymbol{\xi}$ with constant $L_k^y := \frac{1}{\mu_k} + \sigma_{\max}^2(A_1)$, and $\nabla_{\boldsymbol{\xi}} P_k(\boldsymbol{y},\boldsymbol{\xi})$ is Lipschitz continuous in $\boldsymbol{\xi}$ for all fixed $\boldsymbol{y}$ with constant $L_k^\xi := \frac{\gamma}{\mu_k} + \sigma_{\max}^2(A_2)$.

For given $c > 1$ and $\kappa > 0$, it is shown in [19] that the ALCC algorithm, displayed in Fig. 3, can compute an ϵ-optimal and ϵ-feasible solution to (4) within $\mathcal{O}(\log(\epsilon^{-1}))$ ALCC iterations that require at most $\mathcal{O}(\epsilon^{-1}\log(\epsilon^{-1}))$ MAPG iterations. The bottleneck step at each MAPG iteration is the matrix-vector multiplication with $A_1 \in \mathbb{R}^{N^2-N\times N}$, $A_2 \in \mathbb{R}^{N^2-N\times Nn}$, ${A_1}^\mathsf{T}$ and ${A_2}^\mathsf{T}$. Due to specific structures of A_1 and A_2, without forming A_1 and A_2 explicitly, we can compute $A_1 y$ and ${A_1}^\mathsf{T} z$ with $\mathcal{O}(N^2 - N)$ complexity for all y and z; $A_2\xi$ and ${A_2}^\mathsf{T}\omega$ with $\mathcal{O}\big(n(N^2 - N)\big)$ for all ξ and ω. Indeed, neither A_1 nor A_2 is stored in the memory, storing only $\{\boldsymbol{x}_\ell\}_{\ell=1}^N$ is sufficient to be able to compute these matrix-vector multiplications.

Algorithm ALCC ($\boldsymbol{y}_0, \boldsymbol{\xi}_0, \mu_1, \tau_1, \alpha_1^y, \alpha_1^\xi$)
Iteration 0: Take $\boldsymbol{\theta}^0 = \mathbf{0}, k = 1$
Iteration k: ($k \geq 1$)
1) $L_k^y = \frac{1}{\mu_k} + \sigma_{\max}^2(A_1)$, $L_k^\xi = \frac{\gamma}{\mu_k} + \sigma_{\max}^2(A_2)$
2) $\ell_k^{\max} = 4\sqrt{\frac{L_k^y \bar{B}_y^2 + L_k^\xi \bar{B}_\xi^2}{\tau_k}}$
3) $(\boldsymbol{y}_k, \boldsymbol{\xi}_k) = \mathbf{MAPG}\big(P_k, L_k^y, L_k^\xi, \boldsymbol{y}_{k-1}, \boldsymbol{\xi}_{k-1}, \alpha_k^y, \alpha_k^\xi, \ell_k^{\max}\big)$
4) $\boldsymbol{\theta}_{k+1} = \frac{\mu_k}{\mu_{k+1}}\big(A_1\boldsymbol{y}_k + A_2\boldsymbol{\xi}_k - \boldsymbol{\theta}_k\big)_-$
5) $\mu_{k+1} = c\,\mu_k,\ \tau_{k+1} = \tau_k/\big(c\,k^{1+\kappa}\big)^2$
6) $\alpha_{k+1}^y = \alpha_k^y/\big(c\,k^{1+\kappa}\big)^2,\ \alpha_{k+1}^\xi = \alpha_k^\xi/\big(c\,k^{1+\kappa}\big)^2$

Fig. 3: Augmented Lagrangian Algorithm ALCC

Note that at each iteration of ALCC in Step 2) MAPG algorithm is called to inexactly solve (28). Instead of MAPG, one can also use APG in Fig. 1 to inexactly solve (28). Within MAPG algorithm, step sizes taken in each block-coordinate are determined by the block Lipschitz constant, i.e. for y-coordinate the step size is $1/L_k^y$, while it is $1/L_k^\xi$ for the ξ-coordinate. On the other hand, within APG algorithm displayed in Fig. 1, the step sizes taken in each coordinate are equal and determined by the global Lipschitz constant. Thanks to this property of MAPG, we are able to obtain faster convergence in practice in comparison to APG algorithm. When $L_k^y \approx L_k^\xi$, their performance are almost the same; however, when $\max\{L_k^\xi, L_k^y\}/\min\{L_k^\xi, L_k^y\} \gg 1$, since APG uses the global constant L, it takes very tiny steps in one of the block-coordinates.

Algorithm MAPG ($P,\ L^y,\ L^\xi,\ \boldsymbol{y}_0,\ \boldsymbol{\xi}_0,\ \alpha^y,\ \alpha^\xi, \ell^{\max}$)
Iteration 0: Take $\boldsymbol{y}_0^{(1)} = \boldsymbol{y}_1^{(2)} = \boldsymbol{y}_0, \boldsymbol{\xi}_0^{(1)} = \boldsymbol{\xi}_1^{(2)} = \boldsymbol{\xi}_0, t_1 = 1$
Iteration ℓ: ($\ell \geq 1$)
1) $\boldsymbol{y}_\ell{}^{(1)} = \Pi_{\mathcal{Q}_y}\Big(\boldsymbol{y}_\ell^{(2)} - \frac{1}{L^y}\nabla_{\boldsymbol{y}} P(\boldsymbol{y}_\ell^{(2)}, \boldsymbol{\xi}_\ell^{(2)})\Big)$
2) $\boldsymbol{\xi}_\ell{}^{(1)} = \Pi_{\mathcal{Q}_\xi}\Big(\boldsymbol{\xi}_\ell^{(2)} - \frac{1}{L^\xi}\nabla_{\boldsymbol{\xi}} P(\boldsymbol{y}_\ell^{(2)}, \boldsymbol{\xi}_\ell^{(2)})\Big)$
3) $t_{\ell+1} = (1 + \sqrt{1 + 4t_\ell^2})/2$
4) $\boldsymbol{y}_{\ell+1}^{(2)} = \boldsymbol{y}_\ell^{(1)} + \frac{t_\ell - 1}{t_{\ell+1}}\Big(\boldsymbol{y}_\ell^{(1)} - \boldsymbol{y}_{\ell-1}^{(1)}\Big)$
5) $\boldsymbol{\xi}_{\ell+1}^{(2)} = \boldsymbol{\xi}_\ell^{(1)} + \frac{t_\ell - 1}{t_{\ell+1}}\Big(\boldsymbol{\xi}_\ell^{(1)} - \boldsymbol{\xi}_{\ell-1}^{(1)}\Big)$
6) **if** $\|\boldsymbol{y}_\ell^{(1)} - \boldsymbol{y}_\ell^{(2)}\|_2 \leq \alpha^y$ **and** $\|\boldsymbol{\xi}_\ell{}^{(1)} - \boldsymbol{\xi}_\ell^{(2)}\|_2 \leq \alpha^\xi$
7) **return** $(\boldsymbol{y}_\ell^{(1)}, \boldsymbol{\xi}_\ell^{(1)})$
8) **else if** $\ell = \ell^{\max}$
9) **return** $(\boldsymbol{y}_\ell^{(1)}, \boldsymbol{\xi}_\ell^{(1)})$
10) **end if**

Fig. 4: Modified Accelerated Proximal Gradient Algorithm

Convergence and rate result of MAPG follow directly from APG in [14] with the help of following lemma.

Lemma 5: Let $f : \mathbb{R}^m \to \mathbb{R}$ be a convex function, such that $\nabla_{x_1} f(\boldsymbol{x}_1, \boldsymbol{x}_2)$ is Lipschitz continuous with respect to $\boldsymbol{x}_1$ with constant L_1, and $\nabla_{x_2} f(\boldsymbol{x}_1, \boldsymbol{x}_2)$ is Lipschitz continuous in $\boldsymbol{x}_2$ with constant L_2. Then we have

$$f(\boldsymbol{z}_1,\boldsymbol{z}_2) \leq f(\boldsymbol{x}_1,\boldsymbol{x}_2) + L_1\|\boldsymbol{z}_1 - \boldsymbol{x}_1\|_2^2 + L_2\|\boldsymbol{z}_2 - \boldsymbol{x}_2\|_2^2 + \nabla_{\boldsymbol{x}_1} f(\boldsymbol{x}_1,\boldsymbol{x}_2)^\mathsf{T}(\boldsymbol{z}_1 - \boldsymbol{x}_1) + \nabla_{\boldsymbol{x}_2} f(\boldsymbol{x}_1,\boldsymbol{x}_2)^\mathsf{T}(\boldsymbol{z}_2 - \boldsymbol{x}_2).$$

Proof: From Lipschitz continuity of $\nabla_{x_1} f(., x_2)$ for each x_2 and $\nabla_{x_2} f(x_1, .)$ for each x_1, it follows that

$$f(\boldsymbol{y}_1,\boldsymbol{x}_2) \leq f(\boldsymbol{x}_1,\boldsymbol{x}_2) + \nabla_{\boldsymbol{x}_1} f(\boldsymbol{x}_1,\boldsymbol{x}_2)^\mathsf{T}(\boldsymbol{y}_1 - \boldsymbol{x}_1) + \tfrac{L_1}{2}\|\boldsymbol{y}_1 - \boldsymbol{x}_1\|_2^2, \qquad (29)$$
$$f(\boldsymbol{x}_1,\boldsymbol{y}_2) \leq f(\boldsymbol{x}_1,\boldsymbol{x}_2) + \nabla_{\boldsymbol{x}_2} f(\boldsymbol{x}_1,\boldsymbol{x}_2)^\mathsf{T}(\boldsymbol{y}_2 - \boldsymbol{x}_2) + \tfrac{L_2}{2}\|\boldsymbol{y}_2 - \boldsymbol{x}_2\|_2^2. \qquad (30)$$

Multiplying (29) and (30) with $\frac{1}{2}$, and summing them up, gives us

$$\begin{aligned}&\tfrac{1}{2}f(\boldsymbol{y}_1,\boldsymbol{x}_2) + \tfrac{1}{2}f(\boldsymbol{x}_1,\boldsymbol{y}_2)\\ &\leq f(\boldsymbol{x}_1,\boldsymbol{x}_2) + \tfrac{1}{2}\nabla f(\boldsymbol{x}_1,\boldsymbol{x}_2)^\mathsf{T}\begin{pmatrix}\boldsymbol{y}_1 - \boldsymbol{x}_1\\ \boldsymbol{y}_2 - \boldsymbol{x}_2\end{pmatrix}\\ &\quad + \frac{L_1}{4}\|\boldsymbol{y}_1 - \boldsymbol{x}_1\|_2^2 + \frac{L_2}{4}\|\boldsymbol{y}_2 - \boldsymbol{x}_2\|_2^2.\end{aligned}$$

Let $\boldsymbol{z}_1 = (\boldsymbol{x}_1 + \boldsymbol{y}_1)/2$ and $\boldsymbol{z}_2 = (\boldsymbol{x}_2 + \boldsymbol{y}_2)/2$, and by convexity of f, we have

$$f(\boldsymbol{z}_1,\boldsymbol{z}_2) \leq \frac{1}{2}f(\boldsymbol{y}_1,\boldsymbol{x}_2) + \frac{1}{2}f(\boldsymbol{x}_1,\boldsymbol{y}_2).$$

Combining the last two inequality concludes the proof. □

Let $\{\boldsymbol{y}_\ell^{(1)}, \boldsymbol{\xi}_\ell^{(1)}\}_{\ell\in\mathbb{Z}_+}$ be the iterate sequence generated by MAPG algorithm while running on (28) starting from

$(\boldsymbol{y}_{k-1}, \boldsymbol{\xi}_{k-1})$. Using Lemma 5 and adapting the proof of Theorem 4.4 in [14], it can be shown that for all $\ell \geq 1$,

$$0 \leq P_k\left(\boldsymbol{y}_\ell^{(1)}, \boldsymbol{\xi}_\ell^{(1)}\right) - P_k^* \leq \frac{4\left(L_k^y \|\boldsymbol{y}_{k-1} - \boldsymbol{y}_k^*\|_2^2 + L_k^\xi \|\boldsymbol{\xi}_{k-1} - \boldsymbol{\xi}_k^*\|_2^2\right)}{(\ell+1)^2},$$

where $(\boldsymbol{y}_k^*, \boldsymbol{\xi}_k^*)$ is a minimizer of (28). Note that we have $\|\boldsymbol{y}_{k-1} - \boldsymbol{y}_k^*\|_2 \leq 2\bar{B}_y$ and $\|\boldsymbol{\xi}_{k-1} - \boldsymbol{\xi}_k^*\|_2 \leq 2\bar{B}_\xi$. Hence, for all $\ell \geq \ell_k^{\max}$, it is guaranteed that $\left(\boldsymbol{y}_\ell^{(1)}, \boldsymbol{\xi}_\ell^{(1)}\right)$ is τ_k-optimal to (28).

Note that one can also use ALCC, displayed in Fig. 3, to compute the primal iterates $\boldsymbol{\eta}_\ell$ in Step-1 of P-APG in Fig. 2 during the ℓ-th iteration. In particular, at beginning of every P-APG iteration, $\boldsymbol{\eta}_\ell$ can be computed using ALCC to evaluate $\nabla g_\gamma(\boldsymbol{\theta}_\ell^{(2)})$. More importantly, thanks to the separability of regularized (7), one can do this computation in parallel running ALCC on each one of the K processors.

Let $N = K\bar{N}$ such that $N \geq n+1$. Below we consider the bottleneck memory requirement for solving (9) in 4 cases: **a)** P-APG with ALCC computing Step-1 in Fig. 2, **b)** running ALCC *alone* on (9), **c)** P-APG with a primal-dual IPM computing Step-1 in Fig. 2, and **d)** running IPM *alone* on (9). The bottleneck for case **a)** is determined by Step-2 in Fig. 2, due to dual iterates $\boldsymbol{\theta}$ of size $(K^2-K)\bar{N}^2 + K\bar{N}$. Similarly, for case **b)** Step-4 in Fig. 3 requires storing $\boldsymbol{\theta}$ of size $K^2\bar{N}^2$. On the contrary, IPM needs to solve a Newton system in each iteration for both cases **c)** and **d)**. Assuming Cholesky factorization is stored, one needs to keep K lower triangular matrices in memory of size $\bar{N}(n+1)$-by-$\bar{N}(n+1)$ for case **c)**, and to keep 1 lower triangular matrix of size $N(n+1)$-by-$N(n+1)$ for case **c)**. Above discussion is summarized in Table I. Note that running IPM within P-APG reduces the memory requirement significantly by a factor of K in comparison to running IPM alone, e.g. if we partition N observations into $K = 10$ subsets and each subproblem requires 1GB of memory, then running IPM alone requires roughly 100GB, while IPM within P-APG requires only 10GB in total.

TABLE I: Comparison of Memory Usage

	IPM	ALCC
Alone	$\mathcal{O}\big(K^2\bar{N}^2(n+1)^2\big)$	$\mathcal{O}\big(K^2\bar{N}^2\big)$
P-APG	$\mathcal{O}\big(K\bar{N}^2(n+1)^2 + (K^2-K)\bar{N}^2\big)$	$\mathcal{O}\big((K^2-K)\bar{N}^2\big)$

III. Numerical Study

In this section, we provide a comparison in Matlab among the following methods: Sedumi, ALCC, Mosek, P-APG with Sedumi, P-APG with ALCC, and P-APG with Mosek, on problem (9) with increasing dimension. The numerical study is mainly aimed to demonstrate how the performance of each method scales with the dimension of the problem.

First, we start with a small size problem: $n = 5, N = 100$. $\Lambda \in \mathbb{R}^{n\times n}$, $\{\boldsymbol{x}_i\}_{i=1}^N \subset \mathbb{R}^n$ and $\{\epsilon_i\}_{i=1}^N \subset \mathbb{R}$ are generated randomly with all the components being i.i.d. with $\mathcal{N}(0,1)$, and $\bar{y}_i$ are generated according to (1), where $f_0(\boldsymbol{x}) = \frac{1}{2}\boldsymbol{x}^\mathsf{T} Q \boldsymbol{x}$, and $Q = \Lambda^\mathsf{T}\Lambda$. We compare the quality of the solutions computed by P-APG and dual gradient ascent (as the dual function g_γ in (16) is differentiable). In order to compute dual gradient, ∇g_γ, one needs to solve K quadratic subproblems. To exploit this parallel structure, we partition the data into two sets, i.e. $K = 2$. Within both the dual gradient ascent and P-APG, we called ALCC to compute the dual gradients via solving K QP subproblems. Since we allow violations for the relaxed constraints, "duality gap" in the paper is defined as $\boldsymbol{\theta}_k^\mathsf{T} C \boldsymbol{\eta}_k$ at k^{th} iteration. Fig. 5 represents how the duality gap of both methods changes at each iteration. In order to better understand the behavior of P-APG, we report in Fig. 6 the duality gap of P-APG in a smaller scale. Fig. 7 reports the infeasibility of iterates, i.e. $\left\|\left(A_1\boldsymbol{y}_k + A_2\boldsymbol{\xi}_k\right)_-\right\|_2$.

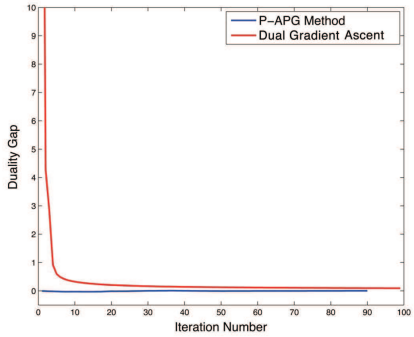

Fig. 5: Duality Gap for P-APG and Dual Gradient Ascent

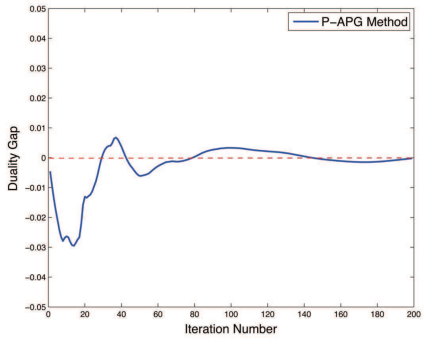

Fig. 6: Duality Gap for P-APG Method

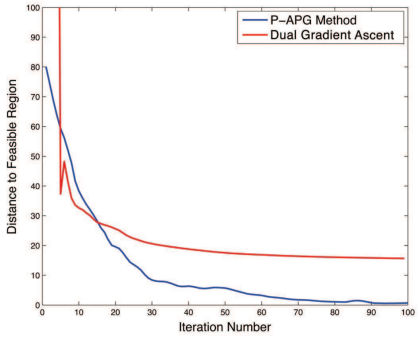

Fig. 7: Distance to Feasible Region for P-APG and Dual Gradient Ascent

TABLE II: Comparison with test function $\exp(\boldsymbol{p}^\mathsf{T}\boldsymbol{x})$

N	Solver	CPU	W.T.	$\frac{1}{2}\|\boldsymbol{y}-\bar{\boldsymbol{y}}\|_2^2$	Gap	Infeas.
200	Sedumi	2.69	2.69	1.16E-05	0	0
	ALCC	1.86	1.86	1.16E-05	-1.19E-07	9.6E-02
	Mosek	0.74	0.74	1.16E-05	2.68E-08	0
	PAPG(Sedumi)	29.64	14.84	1.17E-05	1.31E-07	9.9E-02
	PAPG(ALCC)	9.93	4.98	1.19E-05	2.93E-08	9.9E-02
	PAPG(Mosek)	3.77	1.91	1.17E-05	9.85E-08	9.4E-02
400	Sedumi	O.M.	O.M.	O.M.	O.M.	O.M.
	ALCC	14.74	14.74	6.01E-05	-1.15E-07	9.8E-02
	Mosek	O.M.	O.M.	O.M.	O.M.	O.M.
	PAPG(Sedumi)	120.48	30.28	6.00E-05	4.52E-09	9.7E-02
	PAPG(ALCC)	35.83	9.11	6.11E-05	-2.87E-08	9.8E-02
	PAPG(Mosek)	15.73	4.12	6.00E-05	4.97E-09	9.7E-02
800	Sedumi	O.M.	O.M.	O.M.	O.M.	O.M.
	ALCC	93.57	93.57	2.02E-04	-8.50E-08	9.9E-02
	Mosek	O.M.	O.M.	O.M.	O.M.	O.M.
	PAPG(Sedumi)	146	19	2.46E-04	2.41E-08	9.9E-02
	PAPG(ALCC)	118.54	15.77	2.10E-04	7.01E-08	9.8E-02
	PAPG(Mosek)	52.43	7.52	2.05E-04	6.49E-08	9.7E-02
1600	Sedumi	O.M.	O.M.	O.M.	O.M.	O.M.
	ALCC	N/A	N/A	N/A	N/A	N/A
	Mosek	O.M.	O.M.	O.M.	O.M.	O.M.
	PAPG(Sedumi)	N/A	N/A	N/A	N/A	N/A
	PAPG(ALCC)	323.68	23.94	1.97E-03	5.85E-09	9.9E-02
	PAPG(Mosek)	204.56	17.14	1.85E-03	-6.72E-10	9.9E-02

TABLE III: Comparison with test function $\frac{1}{2}\boldsymbol{x}^\mathsf{T}Q\boldsymbol{x}$

N	Solver	CPU	W.T.	$\frac{1}{2}\|\boldsymbol{y}-\bar{\boldsymbol{y}}\|_2^2$	Gap	Infeas.
200	Sedumi	2.88	2.88	1.25E-04	0	0
	ALCC	3.58	3.58	1.25E-04	-3.41E-08	9.8E-03
	Mosek	0.65	0.65	1.25E-04	2.11E-08	0
	PAPG(Sedumi)	16.7	8.41	1.29E-04	4.48E-06	6.5E-02
	PAPG(ALCC)	12.5	6.3	1.25E-04	9.49E-08	9.8E-02
	PAPG(Mosek)	5.57	2.8	1.26E-04	1.45E-07	8.6E-02
400	Sedumi	O.M.	O.M.	O.M.	O.M.	O.M.
	ALCC	33.3	33.3	1.02E-03	-6.84E-08	1.4E-02
	Mosek	O.M.	O.M.	O.M.	O.M.	O.M.
	PAPG(Sedumi)	164	41.2	1.02E-03	-8.62E-08	9.5E-02
	PAPG(ALCC)	63.2	15.7	1.01E-03	1.53E-07	9.7E-02
	PAPG(Mosek)	29.01	7.43	1.02E-03	-8.74E-08	9.5E-02
800	Sedumi	O.M.	O.M.	O.M.	O.M.	O.M.
	ALCC	140	140	4.03E-02	-2.94E-07	9.9E-02
	Mosek	O.M.	O.M.	O.M.	O.M.	O.M.
	PAPG(Sedumi)	303.33	39.87	4.04E-03	-1.95E-07	9.9E-02
	PAPG(ALCC)	206	23.9	4.04E-03	-2.03E-07	9.9E-02
	PAPG(Mosek)	100.32	14.34	4.04E-03	-1.95E-07	9.9E-02
1600	Sedumi	O.M.	O.M.	O.M.	O.M.	O.M.
	ALCC	N/A	N/A	N/A	N/A	N/A
	Mosek	O.M.	O.M.	O.M.	O.M.	O.M.
	PAPG(Sedumi)	N/A	N/A	N/A	N/A	N/A
	PAPG(ALCC)	480	29.8	5.28E-03	1.10E-07	9.9E-02
	PAPG(Mosek)	273.93	21.47	5.23E-03	5.59E-08	9.8E-02

A primal-dual iterate $(\boldsymbol{\eta}, \boldsymbol{\theta})$ is optimal if the duality gap and infeasibility are both zero. As the feasibility happens in the limit, the duality gap in Fig. 6 can go below the red line, which can be explained by the infeasibility of iterates. Therefore, observing a decrease in duality gap only tells one part of the story; without convergence to feasibility, it is not valuable alone as a measure. As shown in the Fig. 5 and Fig. 6, the duality gap converges quickly to zero for both methods. On the other hand, as shown in Fig. 7, constraint violation for P-APG iterates decreases to 0 much faster than it does for the dual gradient ascent iterates. Hence, P-APG iterate sequence converges to the unique optimal solution considerably faster.

The larger scale problems are carried out on a single node at a research computing cluster. The node is composed of one 16-core processor sharing 32GB. For P-AGPG numerical tests, in each job submitted to the computing cluster, an instance of (9) is solved using P-APG on the node such that each subproblem is computed on a different core. The dimension of variables $n = 80$ and the number of observations $N = 200, 400, 800, 1600$. We partition the set of observations into K subsets. Each one of them consists of 100 points. So, $K = 2, 4, 8, 16$ for $N = 200, 400, 800, 1600$, respectively. In all the tables, *N/A* means that the wall clock time exceeded 2 hours for the job, and *O.M.* means the algorithm in focus runs out of memory. Also CPU denotes the CPU run time in *minutes*; W.T. stands for wall-clock time in *minutes*. Since the number of constraints increases at the rate of $\mathcal{O}(N^2)$, as the size of problem increases in N, we reported the normalized infeasibility and normalized duality gap, which are $\|(A_1\boldsymbol{y} + A_2\boldsymbol{\xi})_-\|_2/\sqrt{N^2-N}$ and $\boldsymbol{\theta}_k^\mathsf{T} C\boldsymbol{\eta}_k/(N^2-N)$, respectively. We report numerical results for the following test functions: $f_0(\boldsymbol{x}) = \frac{1}{2}\boldsymbol{x}^\mathsf{T}Q\boldsymbol{x}$, $f_0(\boldsymbol{x}) = \exp(\boldsymbol{p}^\mathsf{T}\boldsymbol{x})$, where Q is generated as discussed before, and $\boldsymbol{p} \in \mathbb{R}^n$ is generated using uniform distribution.

TABLE IV: Replications with test function $\exp(\boldsymbol{p}^\mathsf{T}\boldsymbol{x})$

Solver	Rep.	CPU	W.T.	$\frac{1}{2}\|\boldsymbol{y}-\bar{\boldsymbol{y}}\|_2^2$	Gap	Infeas.
PAPG(ALCC)	1	118.54	15.77	2.10E-04	7.01E-08	9.8E-02
	2	131.93	17.56	3.90E-04	5.52E-09	9.9E-02
	3	136.50	18.01	4.69E-04	-3.87E-09	9.8E-02
	4	126.43	16.74	2.91E-04	-2.53E-09	9.9E-02
	5	144.31	18.98	5.35E-04	-7.62E-08	9.8E-02
PAPG(Mosek)	1	52.43	7.52	2.05E-04	6.49E-08	9.7E-02
	2	57.33	8.12	3.77E-04	9.39E-09	9.9E-02
	3	61.53	8.64	4.54E-04	-5.30E-09	9.9E-02
	4	55.86	8.00	2.84E-04	-7.48E-09	9.9E-02
	5	65.04	9.15	5.15E-04	-7.74E-08	9.8E-02

TABLE V: Replications with test function $\frac{1}{2}\boldsymbol{x}^\mathsf{T}Q\boldsymbol{x}$

Solver	Rep.	CPU	W.T.	$\frac{1}{2}\|\boldsymbol{y}-\bar{\boldsymbol{y}}\|_2^2$	Gap	Infeas.
PAPG(ALCC)	1	206.00	23.90	4.04E-03	-2.03E-07	9.9E-02
	2	213.27	27.63	1.00E-03	-9.04E-08	9.6E-02
	3	211.18	27.37	1.11E-03	-1.09E-07	9.9E-02
	4	178.77	23.41	7.29E-04	-6.20E-08	9.9E-02
	5	200.62	26.14	1.27E-03	-1.05E-07	9.6E-02
PAPG(Mosek)	1	100.32	14.34	4.04E-03	-1.95E-05	9.9E-02
	2	79.27	10.87	9.88E-04	-1.04E-07	9.9E-02
	3	83.11	11.43	1.09E-03	-1.19E-07	9.9E-02
	4	68.90	9.66	7.16E-04	-2.30E-08	9.9E-02
	5	79.55	10.99	1.24E-03	-1.19E-07	9.8E-02

All the algorithms are terminated either when they compute an iterate with normalized infeasibility and normalized duality gap are less than 1E-01 and 1E-06, respectively, or at the end of 2 hours. The numerical results reported in Table II and III show that P-APG solution is very close to the real optimal solution of (9). Note that ALCC fails to terminate within in 2 hours when $N = 1600$; and interior point methods fail to run anything beyond $N = 200$ due $\mathcal{O}(N^2n^2)$ memory requirement. Moreover, in order to test the robustness of P-APG, we solved 5 random instances when $N = 800$, of which results are reported in Table IV and Table V. Numerical results show that advantages of P-APG over running IPM or ALCC alone on (9) become more and more evident as the dimension of the problem increases.

IV. Conclusion

In this paper, we proposed P-APG method to efficiently compute the least squares estimator for large scale convex regression problems. By relaxing constraints partially, we obtained the separability on the corresponding Lagrangian dual problem. Using Tikhonov regularization, we ensured the feasibility of iterates in the limit, and we provided error bounds on 1) the distance between the inexact solution to the regularized problem and the optimal solution to the original problem, 2) the constraint violation of the regularized solution. The comparison in the numerical section demonstrates the efficiency of P-APG method on memory usage compared to IPM. Furthermore, the extended random tests show the stability of P-APG method. Due to limited space, we could not include computational results on real-life data; but they will be made available online at authors' webpage.

References

[1] M. Mousavi and P. Glynn. Shape-constrained estimation of value functions. *preprint available at arXiv:1312.7035*, 2013.

[2] R. Meyer and J. Pratt. The consistent assessment and fairing of preference functions. *Systems Science and Cybernetics, IEEE Transactions on*, 4(3):270–278, 1968.

[3] H. Chen and D. Yao. *Fundamentals of queueing networks: Performance, asymptotics, and optimization*, volume 46 of *Stochastic Modelling and Applied Probability*. Springer, 2001.

[4] E. Lim and P. Glynn. Consistency of multidimensional convex regression. *Operations Research*, 60(1):196–208, 2012.

[5] P. Groeneboom, G. Jongbloed, and J. Wellner. Estimation of a convex function: Characterizations and asymptotic theory. *Annals of Statistics*, 29(6):1653–1698, 2001.

[6] E. Seijo and B. Sen. Nonparametric least squares estimation of a multivariate convex regression function. *The Annals of Statistics*, 39(3):1633–1657, 2011.

[7] L. Hannah and D. Dunson. Approximate dynamic programming for storage problems. In *Proceedings of the 28th International Conference on Machine Learning (ICML-11)*, pages 337–344, 2011.

[8] L. Hannah and D. Dunson. Multivariate convex regression with adaptive partitioning. *The Journal of Machine Learning Research*, 14(1):3261–3294, 2013.

[9] S. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press, New York, NY, USA, 2004.

[10] J. Nocedal and S. J. Wright. *Numerical Optimization*. Springer, New York, 2nd edition, 2006.

[11] Y. Nesterov. *Introductory lectures on convex optimization*, volume 87 of *Applied Optimization*. Kluwer Academic Publishers, Boston, MA, 2004.

[12] H. Engl, K. Kunisch, and A. Neubauer. Convergence rates for Tikhonov regularisation of non-linear ill-posed problems. *Inverse problems*, 5(4):523, 1989.

[13] D. P. Bertsekas. *Nonlinear Programming*. Athena Scientific, 1999.

[14] A. Beck and M. Teboulle. A fast iterative shrinkage-thresholding algorithm for linear inverse problems. *SIAM J. Img. Sci.*, 2(1):183–202, March 2009.

[15] P. Tseng. On accelerated proximal gradient methods for convex-concave optimization. Preprint available at http://www.eecs.berkeley.edu/˜brecht/eecs227cdocs/tseng.pdf, 2008.

[16] Y. Nesterov. Smooth minimization of nonsmooth functions. *Mathematical Programming, Series A*, 103:127–152, 2005.

[17] Y. Nesterov. Excessive gap technique in nonsmooth convex minimization. *SIAM Journal on Optimization*, 16(1):235–249, 2005.

[18] N. D. Yen. Lipschitz continuity of solutions of variational inequalities with a parametric polyhedral constraint. *Mathematics of Operations Research*, 20(3):695–708, 1995.

[19] N. S. Aybat and G. Iyengar. An augmented lagrangian method for conic convex programming. preprint, arXiv:1302.6322 [math.OC], 2013.